AFRL-RI-RS-TR-2011-164

CORNELL WORKSHOP SERIES ON INFORMATION MANAGEMENT IN SERVICE ORIENTED APPLICATIONS

CORNELL UNIVERSITY, INC

JUNE 2011

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

STINFO COPY

AIR FORCE RESEARCH LABORATORY
INFORMATION DIRECTORATE

■ **AIR FORCE MATERIEL COMMAND** ■**UNITED STATES AIR FORCE** ■ **ROME, NY 13441**

NOTICE AND SIGNATURE PAGE

Using Government drawings, specifications, or other data included in this document for any purpose other than Government procurement does not in any way obligate the U.S. Government. The fact that the Government formulated or supplied the drawings, specifications, or other data does not license the holder or any other person or corporation; or convey any rights or permission to manufacture, use, or sell any patented invention that may relate to them.

This report was cleared for public release by the 88th ABW, Wright-Patterson AFB Public Affairs Office and is available to the general public, including foreign nationals. Copies may be obtained from the Defense Technical Information Center (DTIC) (http://www.dtic.mil).

AFRL-RI-RS-TR-2011-164 HAS BEEN REVIEWED AND IS APPROVED FOR PUBLICATION IN ACCORDANCE WITH ASSIGNED DISTRIBUTION STATEMENT.

FOR THE DIRECTOR:

/s/
RICHARD FEDORS
Work Unit Manager

/s/
JULIE BRICHACEK, Chief
Information Systems Division
Information Directorate

This report is published in the interest of scientific and technical information exchange, and its publication does not constitute the Government's approval or disapproval of its ideas or findings.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE	3. DATES COVERED *(From - To)*
JUN 2011	Final Technical Report	AUG 2009 – JAN 2011

4. TITLE AND SUBTITLE
CORNELL WORKSHOP SERIES ON INFORMATION MANAGEMENT IN SERVICE ORIENTED APPLICATIONS

5a. CONTRACT NUMBER FA8750-09-1-0209
5b. GRANT NUMBER N/A
5c. PROGRAM ELEMENT NUMBER 62702F

6. AUTHOR(S)
Kenneth Birman, Robert van Renesse, Daniel Freedman

5d. PROJECT NUMBER IMCT
5e. TASK NUMBER CL
5f. WORK UNIT NUMBER US

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Cornell University, Inc
Office of Sponsored Programs
373 Pine Tree Rd
Ithaca, NY 14850-2820

8. PERFORMING ORGANIZATION REPORT NUMBER
N/A

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Air Force Research Laboratory/RISF
525 Brooks Road
Rome NY 13441-4505

10. SPONSOR/MONITOR'S ACRONYM(S)
AFRL/RI

11. SPONSORING/MONITORING AGENCY REPORT NUMBER
AFRL-RI-RS-TR-2011-164

12. DISTRIBUTION AVAILABILITY STATEMENT
Approved for Public Release; Distribution Unlimited. PA# 88ABW-2011-1996
Date Cleared: 5 APR 2011

13. SUPPLEMENTARY NOTES

14. ABSTRACT
The Cornell University Distributed Systems research group conducted a lecture series at the Air Force Research Laboratory Information Directorate (AFRL/RI) in Rome, NY. A total of eight half-day workshops on technologies for building robust cloud computing solutions were held at Rome Research Site. During these workshops, AF Research Laboratory and Cornell University participants, as well as attending vendors, engaged in a broad range of exchanges that helped prepare all involved in the more difficult issues surrounding Information Management. These issues needed to be addressed in order to improve mission planning and execution.

15. SUBJECT TERMS
Workshop Series, Information Management, Service Oriented Applications, Cloud Computing

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			RICHARD FEDORS
U	U	U	UU	18	19b. TELEPHONE NUMBER (*Include area code*) N/A

TABLE OF CONTENTS

1.0 SUMMARY

The objective of this effort was to discover, share, and analyze new Information Management concepts that could provide accessible-tailored-actionable information to future warfighters. Additional benefits from this laboratory/university collaboration are to foster stronger military-academic research, aid recruitment of top university talent, and educate the AFRL workforce on the latest advances in the Information Management field. Cornell University's Distributed Systems research group conducted a lecture series at the Air Force Research Laboratory Information Directorate (AFRL/RI) in Rome, NY. A total of eight half-day workshops on information management technologies such as building robust cloud computing solutions; integrating live information objects to provide nimble support tools; and exploring information management challenges in wide-area networks were held at Rome Research Site. During these workshops, AF Research Laboratory and Cornell University participants, as well as attending vendors, engaged in a broad range of exchanges that helped prepare all involved for tackling the difficult issues surrounding military information management. These issues must be addressed in order to improve mission planning and execution in contested, highly dynamic environments.

2.0 INTRODUCTION

There is a growing opportunity to use Service-Oriented Collaboration Applications in ways that can enable more effective search and rescue after a disaster, support a more nimble information-enabled military, and make possible coordinated mission planning and execution even in situations where joint forces or coalition participants create information-sharing obstacles and security challenges. Collaboration applications will need to combine two types of content: traditional web service hosted content, such as data from data bases, image repositories, patient records, and weather prediction systems, with a variety of collaboration features, such as chat windows, white boards, peer-to-peer video and other media streams, and replication/coordination mechanisms.

These applications must also work within Web Services (including the global information grid or GIG) standards, and run over standard Internet infrastructure. In effect, they must leverage cloud computing platforms constructed using scalable data center technologies. Because the U.S. Air Force has selected certain platform technologies for these purposes (Windows system 7 on the desktop and either Windows Enterprise Server of Linux Red Hat on servers), these specific platforms must be employed when possible, and must be used in as standard a manner as possible. Our work shop series will focus on the challenges created by these mixed goals: those of creating powerful new tools for information-informed military missions that run over standard infrastructure, and as much as possible build upon unmodified commodity platforms designated by the Air Force as preferred platforms.

3.0 METHODS

As part of AFRL grant FA8750-09-1-0209, "Workshop Series on Information Management in Service Oriented Applications," the Cornell University Distributed Systems research group has run a lecture series at the Air Force Research Laboratory Information Directorate (AFRL/RI) in Rome, NY. We have held a total of eight half-day workshops on technologies for building robust cloud computing solutions. The first two of these were held in 2009; the remaining ones in 2010.

During these workshops, AF Laboratory and Cornell University participants, as well as attending vendors, engaged in a broad range of exchanges that helped prepare all involved in the more difficult issues surrounding Information Management. These issues needed to be addressed in order to improve mission planning and execution. Our experience is that this has been highly successful. The Cornell participants presented reviews of current technology and research directions, while workshop attendants had ample opportunity to discuss the materials, and to provide input on the direction in which future workshops should go. This interaction was educational to all participants and ensured that the workshops achieved their goals of informing new development of applications that support a more nimble information-enabled military.

Each workshop was accompanied by an extensive collection of PowerPoint slides that were distributed to all participants. Those materials were also made available to members of the Rome Research Site technical staff who were unable to attend the workshop sessions.

4.0 RESULTS

4.1 Cloud Computing I

The first workshop was held September 16, 2009, in which Ken Birman gave an overview of the Cloud Computing and Web Services paradigms. Cloud computing has swept the industry, and the web services community is rapidly adopting the associated technologies, standards and tools. While cloud computing is certainly aligned with military GIG standards, it poses important new obstacles, challenges and opportunities. Expanding on this topic, Dr. Birman surveyed web services and information management from a cloud perspective, focusing on technology support for building services that can handle heavy request rates from large numbers of client platforms. Building such services entails replicating the underlying databases and files, typically using an enterprise service bus (ESB) to relay updates or, in the case of databases, to replicate entire transactions. Neither scheme works as well as one would wish. Of course the server side isn't the whole story. This presentation also touched on the map-reduce paradigm and on data center file systems. The focus was on obstacles and challenges, not products. For example, the CTO of the Air Force and others have pointed to scalability and performance issues with many ESB products, making data replication surprisingly slow in some data center settings. CTOs of many large companies, such as eBay, Amazon, Microsoft and even IBM are finding that transactional mechanisms (atomicity, consistency, isolation, durability or ACID properties) and other forms of consistency can cause performance problems, and are arguing for weaker forms of consistency. Building scalable services, for these and other reasons, can be surprisingly hard, though there may be ways to reintroduce technology support for transactions and consistent replication in ways that can avoid the scalability and performance issues these CTOs worry about.

In the next lecture, Robbert van Renesse considered the challenges associated with protecting data centers against themselves, and against attack. The former problem can occur because of the trend towards weaker forms of consistency in modern cloud settings. As we give up ACID guarantees and other forms of strong consistency, services can more easily fall behind and this might not be immediately obvious. Services can also be implemented with buggy software that sometimes contaminates data or data structures in harmful ways. And then with the tremendous volume of attacks seen both in the public Internet and even in military settings, one has to worry about viruses that could compromise some of the components of a service. There has been a flurry of recent work on protecting services against these kinds of faults: a spectrum that ranges from crash failures to data corruption to outright "Byzantine" behavior. But how realistic is it to consider using such solutions in mission-critical, performance-critical military data centers? Recent progress in this area was reviewed, before looking at some example scenarios, asking how one might introduce stronger guarantees into military information management services and platforms. There are risks today, as well as options for attacking those risks, and the potential for substantial progress exists along the road beyond the immediate future.

In the final lecture of the first workshop, Dan Freedman examined the networking fabric that supports inter-data center communication. There isn't just one data center; modern cloud systems, including military web services, are invariably federated over networks to create integrated solutions that often span a huge geographic range. Yet precisely because these applications have components at multiple locations, under multiple administrators, and depend so crucially on communication, they are often problematic to implement, performance-tune, and operate.

Earlier lectures in this mini-course looked at widely deployed frameworks, such as Enterprise Service Buses (ESB), for constructing cloud-computing solutions at a high level. However, the underlying communications fabrics that transport data throughout the cloud can significantly impact the performance, scalability, and reliability of the resulting solutions that are layered atop them.

This lecture first reviewed some central concepts of data communication stacks and then investigated the differences among a number of deployed transport fabrics. Means of improving such commodity fabrics were presented and discussed, before examining a case study from our recent research in empirically characterizing high-speed private optical data links. Finally, we discussed implications for improving data transport in the cloud for the forward-deployed expeditionary air force.

4.2 Cloud Computing II

The second workshop, held on 23 November 2009, focused entirely on networks for cloud computing. In the first talk, Hussam Abu-Libdeh and Daniel Freedman examined issues involved with bringing Internet Protocol (IP) Multicast to the Data Center. IP Multicast has many well-known issues that have caused data center operators to minimize or in some cases abandon using it, even though it is a hardware-supported building block of many distributed systems. Characterizing those IP Multicast issues, Dr. Abu-Libdeh went on outlined efforts underway at Cornell to solve them. He discussed work targeting the group scalability issues and rate-limiting issues of IP Multicast. He also described a heuristic they developed for compressing group memberships based on analyzing traces of existing data center applications. Next he covered flow-control and rate-limiting in multi-group settings. Although group communication flow-control was studied extensively two decades ago, most of that work considered a single group. The current work considers rate-limiting in the presence of multiple groups where the behavior of a single group might not be problematic, but where the behavior of the system as a whole is. To address this, they have worked on reactive soft rate-limiting, and plan to do future work on proactive rating-limiting. Dr. Abu-Libdeh believes that this work is an important step towards making IP Multicast usable in data centers.

The second presentation, by Hakim Weatherspoon, was entitled “Commodity Middleboxes Connecting DataCenters Over Vast Geographical Distances.” It started with the observation that security, reliability, and performance are paramount in networks of globally distributed datacenters, especially military networks. Example functionality for these networks includes deep packet inspectors (DPI), wide-area performance enhancement proxies (PEP), protocol accelerators, overlay routers, security appliances, intrusion detection systems (IDS), and network monitors, to name a few. Networks are increasingly dependent upon packet processing network appliances - middleboxes - that reside at the edge of a datacenter network and need to run at wire speed. These ‘appliances’ are key to connecting datacenters over large geographic distances, and enable new network functionality. The typical middlebox is an expensive, high-end piece of equipment running dedicated software, situated at the edge of the datacenter and facing the LAN on one side, and the WAN on the other. This talk discussed how such middleboxes may be built from the same commodity components readily available within the datacenter. It was shown that such performance enhancement proxies can sustain network line speeds (1-10Gbps), and can trivially scale, thus providing a readily available commodity alternative to otherwise proprietary, dedicated equipment. This additional flexibility greatly enhances network functionality.

"Transparent Instrumentation for Reproducible Measurement of 10 Gbps Ethernet" was the next presentation, where Dr. Freedman introduced BiFocals -- transparent instrumentation that enables reproducible optical network measurements and systematically characterizes the source, degree, and propagation of error. BiFocals achieves highly accurate network tomography by leveraging the combination of a precisely calibrated external hardware time-base from an oscilloscope with post-processing software modules. He described the resulting metrological precision, and associated reproducibility and reliability, which was demonstrated empirically by characterizing 10 Gbps Ethernet packet flows in flight across a short fiber-optic link in a laboratory setting. In contrast with many common network tomography methods that use software applications on link endpoints to observe network flow and infer network behavior, BiFocals enables reliable in-situ measurements of packets in flight directly on the wire, without interfering with the system under test. BiFocals thus avoids certain systematic distortions that may otherwise arise, qualitatively altering the distribution of packet timings. His team showed that this approach achieves up to six orders of magnitude improved precision in packet delay measurements over many common tomography methods.

4.3 Live Distributed Objects

The third workshop was held at the Griffiss Institute on March 4th 2010. Cornell University researchers presented a half-day workshop on Live Distributed Objects, a new technology for building collaboration tools on the cloud. Using a drag-and-drop approach, developers can use live objects to create and share applications combining cloud-hosted content; database and enterprise applications data; real world video or radar sensor inputs; and even synthetic sources such as avatars.

Looking at Live Distributed Objects (LDO) from various perspectives, Professor Ken Birman first led a review of the technology as a whole, showing how this Web 2.0 technology for building collaboration tools fits into Information Management platforms and systems, and how it can change the game for information application developers. This technology has also been briefed to three successive CIOs of the Air Force, to senior staff at Microsoft, and to the CTO of Intel. Discussions with these individuals indicate that LDO could be a good match for their organizations' need. For example, in tactical Air Force settings, it can be quite hard to use existing "wired" mashup technologies, like Google Wave or SilverLight, to support warfighters who are often operating with limited connectivity and who may need data from multiple services or even from coalition partners with different security policies than our own. With LDO such problems can be solved; whereas industry standards often scale poorly or shut down if reach-back links are not available.

Dr. Krzys Ostrowski next presented a more detailed technical perspective on Live Distributed Objects. He is the lead developer for the platform, and explained how LDO helps to reduce the inherent complexity in constructing complex distributed applications. He reviewed the disadvantages of the low-level endpoint-centric approach underpinning existing platforms, which motivated a high-level design methodology based on the new abstractions: distributed objects and information flows. He explained how the new abstractions may be used to represent hierarchical systems; how one can translate such high-level representations into code; and how one can reason about consistency in the resulting architectures.

The third part of the workshop was hands-on and practical, with Dr. Freedman showing how his team of Cornell MEng students -- skilled developers but with little special knowledge of distributed protocols -- created a potentially useful application of LDO for a tactical response system aimed at countering a hypothetical terrorism event in the New York City region. A second such system under development will illustrate the use of LDO for medical collaboration. Dr. Freedman described the existing solution, demonstrating the steps required to extend it, and talked about some of the performance and scalability thinking that the developer needs to engage in when using this new kind of mashup technology. This work is relevant to other kinds of cloud platforms, including Google Wave and Microsoft Silverlight. Attendees learned about the LDO platform, were introduced to downloading and using it, and saw how it might be applied in other technology areas such as exploiting multicore computing capabilities. LDO is easy to extend and treats everything, even data replication protocols, as a kind of object. Thus if some new information source or replication protocol is needed, one can easily extend the platform to support it.

4.4 Challenges for Tactical Network Applications

Initially titled "Towards a Scientific Basis for Tactical Network Applications," the fourth workshop was held May 11th 2010 at the AFRL Rome Research site. Dr. Ken Birman and Dr. Robbert van Renesse from Cornell spoke on cloud computing, underlying web service protocol standards, and the challenges when those standards need to cope with mobility, disconnection, or non-wired network architectures.

The dramatic evolution of web-services computing has transformed the landscape. Trade publications suggest that HTTP is the new 'Internet protocol' and, of course, client-server based computing through the web services paradigm underlies much of what has come to be called cloud computing in the popular press. Yet these trends have left mobile users at a disadvantage and frequently disconnected users doubly so. Two examples of this for a tactical setting are (1) securely sharing video intelligence, and (2) coordinating to carry out a simple mission that involves access to evolving information from a shared database hosted on the cloud system.

Professor Birman addressed the issue first from a network "engineering" perspective, and then as a more principled "theory" perspective. A small set of foundational building blocks was presented, in order to identify basic questions that need to be solved in order to support the kinds of functionality familiar from information applications hosted in the cloud, but now running with many clients in a complex tactical networking setting.

The second lecture, given by Dr. van Renesse, looked at these building blocks carefully. He reviewed the relevant theory, which limits what can and cannot be accomplished and offers insight into the costs, and then developed solutions for our examples. The discussion then turned to how those solutions might be implemented while also remaining as true to the web services model as possible. The reality is that tactical networking has lagged behind cloud computing in many ways. The wrap-up session led by both speakers pulled some of the insights gained into a small set of research challenges. Progress on these challenge topics will be the key to enabling much more "nimble" applications that bring information-enabled support to the mobile user operating in a difficult environment.

4.5 Advances in Image-Based Applications and Robotic Vision

At the fifth workshop on July 8, 2010 at the Griffiss Institute, Cornell University professors Ashutosh Saxena and Noah Snavely presented their research on advanced image-based applications.

Professor Saxena explained development and use of the Make3D image-processing algorithm in a talk entitled "Single Image Depth Perception and its applications to Robotics." He considered the challenge -- based on the loss of depth information -- of converting standard digital pictures into three-dimensional (3D) models. This is a difficult problem, since a still image is formed by a projection of the 3D scene onto two dimensions, thus losing the depth information. This method uses a supervised learning approach to model scene depth as a function of image features, and shows that even on unstructured scenes in a variety of environments, it is frequently possible to recover accurate 3D models. He then described recent demonstrations of this new image processing tool in (a) obstacle avoidance for autonomously driving a small electric car at high speeds through cluttered environments and flying indoor helicopters; and (b) robot manipulation, where learning algorithms for grasping novel objects are being developed. This enables a robot to perform tasks such as open new doors, clear cluttered tables, and unload items from a dishwasher. The ability to recover accurate 3D models from images has wide application for autonomous land and air vehicles operating in cluttered environments.

Professor Snavely's talk, "Building Rome in a Day: City-Scale 3D Reconstruction from Image Collections," addresses the problem of organizing and using vast amounts of visual data. We live in a world where visual data is becoming ubiquitous, but increasingly difficult to manage. The Internet, in particular, has become an unprecedented source of visual information about our world, with billions of images on photo sharing and social networking sites like Flickr and Facebook. For instance, we can easily find millions of photos of cities: a Flickr search for "Rome" returns over two million photos, for "NYC" over five million, and for "London" over eight million. This talk described new computer vision techniques that have been developed for taking large, unstructured image collections such as these, and automatically reconstructing city-scale 3D models of the underlying scenes. New computer graphics and visualization techniques for exploring and navigating these reconstructed scenes were also discussed. These techniques have wide applicability, and can be used with any large collection of photographs gathered from heterogeneous sources.

4.6 Networking and Computation in the Next Frontier

Technical advances in recent decades have transformed the landscape of networked computing systems. New networking technologies such as wireless ad hoc networks have enabled computing devices to expand deeply into the physical world. The hardware platforms have also evolved to enable extremely low-power operations with strong security guarantees, even in physically exposed environments. In the sixth workshop "Networking and Computation in the Next Frontier," held August 6th 2010 at Rome Research site, Cornell Professors Zygmunt Haas and Edward Suh surveyed emerging networking and hardware technologies, and discussed how the new technologies may impact future military network system designs.

In the first part, Professor Zygmunt J. Haas described three new networking technologies: the Ad Hoc Networks, the Sensor Networks, and the Delay-Tolerant Networks. He discussed how these technologies evolved, how they differ one from the other, and assessed their technology-transfer potential. The basics of Mobile Ad-Hoc Networking (MANET), including issues such as routing, MAC, MANET security, MANET mobility, topology control, and cooperative (MANET) networks were covered as well. Much of the talk concentrated on a number of selected theoretical and practical challenges in making these technologies attractive for potential military and commercial exploitation, following Professor Haas' emphasis on tactical networking. Topics related to these three technologies such as cooperative networking, network-coding, stochastic routing, statistical quality-of-service, topology control, scalability, security support, transmission scheduling, and mobility modeling were also reviewed.

A sample application Dr. Haas introduced was his animal habitat monitoring project, which applies some of the concepts from ad-hoc and sensor networking to monitor the behavior of ocean life. That project uses a novel concept called the Shared Wireless Infostation Model (SWIM). SWIM is an example of Delay-Tolerant Networks that improve the overall capacity-delay tradeoff and optimizes the energy-delay tradeoff.

Professor G. Edward Suh spoke on new hardware technologies that can improve the security and reliability of future systems. This included hardware authentication technologies as well as state-of-the-art extensible hardware technologies. Along with the focus on security, he discussed energy efficiency for hardware platforms. Embedded and mobile platforms often need to operate on a battery or harvest energy from an environment such as light, electromagnetic fields, and vibration. Dr. Suh provided a survey of the current hardware capabilities in such energy-constrained environments and how emerging technologies in non-volatile memory may enable reliable and energy-efficient operation under energy constraints.

The embedded and mobile platforms also face unique security challenges because they are often disseminated, unsupervised, and physically exposed. Professor Suh explained how emerging hardware techniques such as Physical Uncloneable Functions and trusted processor architectures can provide strong security guarantees in confidentiality, integrity, and availability of a system even from adversaries with physical access. The talk will also briefly discuss how these techniques can be applied to other emerging computing platforms such as cloud computing.

4.7 Novel Security for Next-Generation Information Management

On October 22, 2010, we held the seventh workshop in this series, entitled "Novel Security for Next-Generation Information Management," with contributions from Dr. Michael Clarkson, Dr. Willem de Bruijn, and Mr. Jed Liu.

The initial presentation for this workshop was titled "Civitas: Toward a Secure Voting System." Dr. Clarkson described his work on Civitas, an Internet voting system developed at Cornell to provide secure on-line voting. He reviewed the strong, sometimes conflicting security requirements for such voting systems. While voting privacy must be maintained, voters must also be able to verify that the election tally is correct, to ensure that faulty or malicious voting components cannot commit fraud. For electronic voting systems, neither requirement is easy to satisfy individually. Satisfying them simultaneously has proven surprisingly difficult throughout thirty years of research.

Dr. Clarkson discussed how Civitas offers a strong defense against voter coercion and enables verification of election results, and how these features introduce certain tradeoffs. Such security, however is not free, and remains among the open research problems still being addressed.

Following that presentation, Dr. de Bruijn discussed his work on "SafeCard: Defense in Depth in Network Intrusion Prevention." He began with an overview of network intrusion prevention, and how essential it is to containing malware and providing early identification of targeted attacks. High-rate traffic analysis is a hard systems problem, however, that imposes real-time constraints, saturates I/O capacity and involves computationally expensive algorithms. This drove development of ‘SafeCard,’ an intrusion prevention system that integrates pattern matching, flow analysis and polymorphism detection into a multi-stage prevention system, and ‘Streamline,’ a Linux network stack for the rapid development of such systems that minimizes I/O overhead and integrates smart network interface cards (NICs). His talk reviewed common network attacks, presented successful single-purpose defenses, and focused on operating system mechanisms for combining these to offer defense in depth at high rates.

Finally, Mr. Jed Liu presented his work on ‘Fabric,’ a platform that uses higher-level abstractions to build secure federated information systems. Fabric is decentralized, allowing network sites to securely share information and computation resources. Several of its technical innovations address the weaknesses of previous related systems. Both computation and information can move between network nodes to meet security requirements or to reduce bandwidth. Fabric provides a rich Java-like object model. Data resources are labeled with confidentiality and integrity policies enforced through a combination of compile-time and run-time mechanisms. Optimistic, nested transactions provide secure consistency across all objects and nodes. A peer-to-peer dissemination layer provides increased availability and balances load.

Distributed information systems play an increasingly important role in our lives, from finances, to medicine, to our interaction with government agencies. Current practice lacks general, principled techniques for implementing functionality of these systems while also satisfying security and privacy requirements. Mr. Liu explained just such a general method -- the distributed Fabric system -- and detailed the technical innovations it uses to remedy weaknesses of previous related systems. He also presented results from several applications showing that Fabric has a clean and concise programming model, offers good performance, and enforces security.

4.8 Edge and Core Routers

The eighth (and last) workshop, “Edge and Core Routers: A Tutorial” was held the morning of November 19th 2010, and presented by Prof. Ken Birman and Dr. Robbert van Renesse. This workshop reviewed the state of routing, looking both at edge routers (including functionality such as firewalls and network address translation) and core Internet routers (which are ‘hands-off’ and just route as fast as possible). The Border

Gateway Protocol (BGP) was described in some detail, including how BGP works, what issues have emerged over the past decade, how people have responded, and what might happen next. The architecture of routers and switches were presented, as well as the architecture of entire datacenters. Finally, Birman and Van Renesse discussed several research questions in the area of robust internetworking, including:

- The Internet is ossified. Even a single transition to IPv6 is difficult.
    - How can we change the Internet architecture to experiment and deploy novel protocols?
    - What is the best way to provide high assurance in the Internet?
    - Can we improve the mobile experience for roaming users?
    - Can we use formal methods to verify policy compliance?
    - Can we support content-based addressing and reduce duplication of content?
- The control planes of routers and switches have become very complex, and possibly fragile as a result.
    - Can we bring modern virtualization approaches to routers in order to address this?
    - Can we apply formal methods to verify the correct behavior of routers?
    - How do we recover from black-outs, or how do we bootstrap the Internet?
- The core of the Internet is suddenly prone to overload because of the rise in multimedia content while capacity is relatively stagnant because the relatively slow rise in paying customers.
    - How can we provide predictable performance?
    - Can we provide differentiated service without punishing ordinary users?
    - How do we prevent DDoS attacks?
    - How do we prevent spam and botnets?

5.0 CONCLUSION

The "Information Management in Service-Oriented Applications" workshops with AFRL focused on collaborative services. Leading-edge research on topics such as cloud computing, live information objects, scalable web services, fault tolerant storage, high-speed remote backup, and operating system extensions were explored. The interaction with Information Directorate scientists and engineers exploring leading-edge technology is expected to continue generating new ideas for meeting warfighter information needs, as well as speed the transition of the latest technology to the field.

Some take-away messages of these workshops is that currently deployed cloud computing facilities do not provide sufficient reliability for warfighter applications, and that much research is going on, and needed, in order to bridge the gap. This extends from low-level networking needs, both in the wired and wireless realms, to datacenter management, to middleware services, to applications, and to security needs. Fortunately, there is ongoing research in each of these areas, and continually forthcoming results that can contribute to scalable and dependable information management.

6.0 REFERENCES

1. “Cloud Computing I,” “Cloud Computing Scalability,” and “The Wisdom of the Sages,” Dr. Ken Birman, Professor of Computer Science, Cornell University, 16 September 2009.

2. “Data centers under attack!” Dr. Robbert Van Renesse, Principle Research Scientist, Cornell University, 16 September 2009.

3. “Communication in the Cloud,” Dr. Daniel Freedman, Post-Doctoral Research Associate, Cornell University, 16 September 2009.

4. “Towards Making IP Multicast Usable in Today’s Data Centers,” Dr. Hassam Abu-Libdeh, Cornell University, 23 November 2009.

5. "Commodity Middleboxes Connecting DataCenters Over Vast Geographical Distances," Professor Hakim Weatherspoon, Cornell University, 23 November 2009.

6. "Transparent Instrumentation for Reproducible Measurement of 10 Gbps Ethernet," Dr. Daniel Freedman, Post-Doctoral Research Associate, 23 November 2009.

7. “Live Distributed Objects,” Professor Ken Birman, Cornell University, 4 March 2010.

8. “Live Objects Demonstration for a Tactical Response System,” Dr. Daniel Freedman and Dr. Krzysztof Ostrowski, Post-Doctoral Researchers, Cornell University, 4 March 2010.

9. “Methodology for Design and Development of Complex Distributed Applications,” Dr. Krzysztof Ostrowski, Post-Doctoral Research Associate, Cornell University, 4 March 2010.

10. “Responding to the Tactical Networking Challenge,” Professor Ken Birman, Cornell University, 11 May 2010.

11. “Communication Building Blocks for Tactical Network Applications,” Dr. Robbert Van Renesse, Principle Research Scientist, Cornell University, 11 May 2010.

12. “Make 3D: Single Image Depth Perception and Its Applications to Robotics,” Professor Ashutosh Saxena, Cornell University, 8 July 2010.

13. "Building Rome in a Day: City-Scale 3D Reconstruction from Image Collections,” Professor Noah Snavely, Cornell University, 8 July 2010.

14. “Networking and Computation in the Next Frontier,” Professor Zygmunt Haas, Cornell University, 6 August 2010.

15. “Hardware Platforms for the Next Frontier,” Professor G. Edward Suh, Cornell University, 6 August 2010.

16. “Civitas: Toward a Secure Voting System,” Dr. Michael Clarkson, Cornell University, 22 October 2010.

17. “SafeCard: Defense in Depth in Network Intrusion Prevention,” Dr. Willem de Bruijn, Cornell University, 22 October 2010.

18. “Fabric: Higher-Level Abstractions for Building Secure Federated Systems,” Dr. Jed Liu, Cornell University, 22 October 2010.

19. “Network Routers” and “Challenges in Modern Routers,” Professor Ken Birman, Cornell University, 19 November 2010.

20. “Internet Routing,” and “Routing in a Data Center,” Professor Robbert van Renesse, Cornell University, 19 November 2010.

LIST OF ACRONYMS

ACID	Atomicity, Consistency, Isolation, Durability
AFRL	Air Force Research Laboratory
BGP	Border Gateway Protocol
CIO	Chief Information Officer
CTO	Chief Technology Officer
DDoS	Distributed Denial of Service
DPI	Deep Packet Inspectors
ESB	Enterprise Service Bus
GIG	Global Information Grid
HTTP	HyperText Transfer Protocol
IDS	Intrusion Detection Systems
I/O	Input / Output
IP	Internet Protocol
LAN	local area network
LDO	Live Distributed Objects
MAC	Medium Access Control protocol
MANET	Mobile Ad-Hoc Networking
MEng	Master of Engineering
NICs	Network Interface Cards
PEP	Performance Enhancement Proxies
SOA	Service Oriented Applications
SWIM	Shared Wireless Infostation Model
WAN	Wide Area Network